

Thunder's Power

A perspective view of a server room aisle. The aisle is flanked by rows of server racks on both sides. The floor is covered with a grid of perforated metal tiles. In the center of the aisle, there is a large, stylized graphic of a lightning bolt striking down, with a purple and white color scheme. The title "Thunder's Power" is overlaid in large, bold, pink letters across the top of the image.

Delivers Breakthrough Science

Livermore's Thunder supercomputer allows researchers to model systems at scales never before possible.

WHETHER a computer is simulating the aging and performance of a nuclear weapon, the folding of a protein, or the probability of rainfall over a particular mountain range, the necessary calculations can be enormous. With each advancement in system hardware or software, Livermore researchers are ever ready to test the limits of computing power in search of answers to these and other complex problems. To assist them in their research, the Computation Directorate provides supercomputers that often rank high in performance. Today, the Advanced Simulation and Computing (ASC) Program's 360-trillion-floating-point-operations-per-second (teraflops) BlueGene/L supercomputer at Livermore holds the world's number one spot on the LINPACK benchmark, the industry standard used to measure computing speed.

As computing capability increases, Livermore's high-performance computers are assuming an increasingly central role in the Laboratory's research. In the 1990s,

the Accelerated Strategic and Computing Initiative, now called the ASC Program, began using massively parallel computers, with the necessary computational speed and memory capacity to simulate complex multiphysics problems. In the last decade, the supercomputers' supporting role has grown to accurately and quantitatively connect disparate phenomena in basic and applied science research.

To exploit the machines' integrative capability and maximize potential scientific discoveries, physicist Michael McCoy of the Computation Directorate initiated the Multiprogrammatic and Institutional Computing (M&IC) Program in 1996 with institutional support from Jeff Wadsworth, then deputy director for science and technology, and director Bruce Tarter. The program provides Laboratory scientists and engineers access to ASC-class supercomputers for unclassified, mission-related research requiring the capability provided by big computers.

“M&IC provides the Laboratory’s researchers a place to run unclassified calculations that are often as large as those used for the multiscale and multiphenomena calculations associated with the aging and performance of the nation’s weapons stockpile. The simulations we produce differentiate the Laboratory from other institutions,” says McCoy.

Power in Clusters

One advantage of parallel computers is their scalability. Traditionally, parallel computers relied on vendor-integrated systems and software and on proprietary interconnects, where the cost per teraflops improves very slowly. In 2000, McCoy proposed that M&IC switch to Linux cluster technology, which uses large groups

of commodity microprocessors (computer systems that can be manufactured by multiple vendors), third-party interconnects, and open-source Linux operating system software. The first supercomputer assembled using this approach was the Multiprogrammatic Capability Resource (MCR), which includes 2,304 processors capable of 11 teraflops. In 2004, MCR was joined by Thunder, a 23-teraflops cluster of 4,096 processors. At its debut, Thunder was ranked the world’s second fastest supercomputer.

Hal Graboske, then deputy director for science and technology, proposed allocating Thunder as a capability machine by awarding time on it to five Laboratory projects whose calculations required a significant percentage of the system. He dubbed the projects “Thunder Grand Challenges.” These projects ranged from research in climate change to the interface of water and vapor to protein folding. Each Grand Challenge was strategically aligned with a Laboratory mission. Computer scientist Brian Carnes, who now leads M&IC, says, “The Thunder Grand Challenges have made important contributions to the scientific community because the machine’s computational power allows researchers to accurately model systems at a scale never before possible.”

Tracking Climate Change


Similar to the phenomena involved in stockpile stewardship, modeling climate change is a multiphenomena statistical science that requires multiple simulations. These runs take into account atmospheric and oceanic dynamics and physics, topography, vegetation, chemical constituents and aerosols, and the effects from clouds, precipitation, and ice. Climate change research has been focused mostly on global-scale models at resolutions no finer than 100 to 200 kilometers because of computer capability limitations. Doug Rotman, in Livermore’s Carbon Management and Climate Change Program, is leading a Thunder Grand

Challenge to model regional climate change at resolutions of about 10 kilometers. The calculations will allow, for the first time, evaluation of whether changes in certain regional phenomena, such as river flow, are due to anthropogenic (human-induced) factors.

The project team’s strategy has been to perform multicentury, global simulations at 100-kilometer resolution, which is finer than that of any other century-long comprehensive climate run to date. Rotman says, “Comprehensive climate calculations must simulate centuries to millennia in order to properly represent the coupling between the atmosphere and oceans, particularly the deep ocean currents.”

The Community Climate System Model achieves roughly 10 simulated years per computing day using 472 of Thunder’s processors and includes a calculation that simulates a period of more than 1,100 years. Results of the global simulations serve as initial conditions and boundary data for 10-kilometer-scale atmospheric calculations that focus on areas the size of U.S. states. “No previous simulations could provide the resolution or length in time to assess internally driven regional climate variability, which is required to detect climate change at small scales,” says Rotman. The calculations allow researchers to separate human-induced effects from natural variability at subcontinental scales.

The team collaborates with the Scripps Institution of Oceanography of the University of California (UC) at San Diego, the National Center for Atmospheric Research, and the University of Michigan. Researchers at Scripps will use the results of the regional calculations to drive a hydrology model to predict river flows. Comparison of these predictions with observational data will help ascertain the extent of regional climate change. The results will also have important implications for environmental issues such as air and water quality and the management of healthy ecosystems.



(a) Winter precipitation patterns from the Community Climate System Model agree well with (b) observed precipitation from 1971 to 2000, particularly in the western U.S.

Water's Mysteries Uncovered

To run the necessary calculations for large simulations, Livermore researchers often rely on codes that use first principles to characterize material behavior. First-principles molecular dynamics simulations use the laws of quantum mechanics to describe the electrons in a system. The data are then used to compute the interactions between atoms without experimental input. Livermore has developed more than 50 classified and unclassified codes that use first principles.

Even a seemingly basic molecule, such as water, has many unresolved mysteries because researchers have not had the computational power and first-principles codes until now. Yet, understanding water is important for areas ranging from biological and physical sciences to atmospheric and planetary sciences. Thunder's power allows researchers to tailor calculations for individual studies by combining codes. Chemist Chris Mundy and his team have combined codes to resolve long-standing questions about water's behavior. (See *S&TR*, October 2005, pp. 12–18.) They were first to use Monte Carlo methods with first-interaction

potentials to determine the vapor–liquid coexistence curve for water. “You can see water’s vapor–liquid coexistence, and from the simulation, determine the thermodynamic properties,” says Mundy. “This simulation has never before been done directly.” Their results further understanding of chemistry at aqueous interfaces.

The team also has redefined the predicted phase diagram of water in the interior of large planets, confirming a theory that a superionic phase of water exists in the large planets. In this phase, oxygen atoms are fixed, or frozen, while hydrogen atoms move freely. Mundy says, “There has been a great debate about the existence of superionic water. Our calculations reveal that this exotic phase may exist.”



Controversy also exists among experimentalists over the structure of water on Earth due to difficulty in interpreting results such as x-ray data of hydrogen-bond energies at different temperatures. Simulations performed on Thunder in collaboration with Richard Saykally’s research group at UC Berkeley, which is well known for its experimental studies on water, are helping to resolve this controversy. Mundy’s team also collaborates with researchers from the

University of Minnesota, UC Irvine, and the University of Zurich.

One of the team’s next goals is to use Thunder and molecular dynamics simulations to study the biochemical synthesis of uracil, an RNA nucleotide involved in enzyme production. The reaction that occurs in an enzyme is one of the most disputed mechanisms among scientists. The simulations will include a system of more than 30,000 atoms and will be the largest calculation of enzyme catalysis ever conducted. The study may help further understanding of how protein motion affects the efficiency of important enzymatic reactions.

Free from Energy Barriers

As proteins are synthesized, they fold into specific shapes that are related to their intended function. They fold on average in one-thousandth of a second and measure 2 to 3 nanometers long. Because a protein’s structure is tied to its function, Livermore scientists, and many others worldwide, conduct studies of folding processes. Much of this work is influenced by experiments conducted for the Critical Assessment of Structure Prediction biennial event, which is organized by Livermore’s Chemistry, Materials, and



(a) The hydrogen atoms (blue) in water molecules have a tendency to orient themselves into the air at the surface. This orientation may increase reactivity with other types of molecules.
 (b) Simulations of sodium chloride and water molecules show greater interaction between molecules at the surface and a decrease in interaction for molecules in bulk liquid.

Life Sciences Directorate. (See *S&TR*, December 2004, pp. 12–17.)

Predicting the three-dimensional structure of a protein from its amino-acid sequence is extremely complex. However, determining the final structure is only half the problem. The other half is understanding how proteins fold, that is, the sequences of structural changes they undergo before reaching their final native structure.

Scientists estimate humans have between 50,000 and 100,000 different proteins. Many of the protein chains are composed of hundreds of amino acids. Because protein-folding processes are complex, researchers have mostly studied smaller proteins composed of chains of about 50 amino acids. One reason proteins are a challenge to study is that they exist in a watery environment. Researchers must therefore consider the individual atoms

defined by the amino-acid sequences as well as those from the surrounding water molecules. Complicating this scenario is that each atom is dynamic and its state is in constant flux, producing fluctuating energy barriers. A protein must overcome these barriers as it moves into its folded state.

Simulating this ever-changing energy landscape has traditionally been limited to small systems over inadequate time scales. Theoretical physicist Farid Abraham is leading a Thunder Grand Challenge to simulate protein folding over a broad range of temperatures to reveal the thermodynamic pathways.


Abraham and collaborators Jed Pitera and William Swope at IBM Almaden Research Center simulated folding of the 1BBL protein, whose 40 amino acids are part of a larger protein involved in enzymatic activity. Researchers at other

institutions have studied 1BBL, and their experimental results indicate that it folds in a downhill energy process. “Experimentalists can’t measure each wiggle of a protein as it folds,” says Abraham. “However, each of the atoms’ dynamics determines the folding process.”

The team uses a new technique called Replica Exchange Molecular Dynamics (REMD) to simulate different copies (replicas) of the system at the same time but at different temperature values. Each replica evolves independently by molecular dynamics. High-temperature simulation segments facilitate the crossing of the energy barriers. Low-temperature segments explore in detail the conformations present in the minimum energy basins appropriate for the native state of the protein. The result of this swapping between temperatures is that high-temperature replicas help the low-temperature replicas to jump across the energy barriers of the system.

Simulations for the 1BBL protein comprised about 65,000 atoms, most of which were water molecules. The team created 256 replicas of this system and simulated each replica at a different temperature, ranging from 250 to 600 kelvins (body temperature is 300 kelvins). Each replica used four processors running concurrently for more than 2 million central-processing-unit hours.

“Researchers usually use supercomputers to simulate large systems,” says Abraham. “In this case, we’re using a supercomputer to simulate many small systems, while also taking advantage of its power to study a high-temperature state for better sampling of the low-temperature state. There has been an ongoing controversy among experimentalists over the last four years about whether the new downhill paradigm or the old two-state picture describes 1BBL folding. Our results confirm the experimental findings that the protein folds in a downhill energy process.” However, Abraham says that the story is more complicated. “The 1BBL protein does have an initial energy barrier it has to overcome



The Replica Exchange Molecular Dynamics technique simulates protein-folding thermodynamics over a broad range of temperatures. Results for the 1BBL protein show an abrupt transition from the unfolded state to a globular state at a temperature of 380 kelvins.

in order to collapse into a globular package, but then it slithers into its native state through a downhill path, establishing the new folding paradigm.”


The simulations are three times larger and four times longer than previous studies. However, Abraham says it’s not just having access to the best supercomputers in the world that enables the Laboratory’s discoveries. “Our results were possible because of M&IC’s dedication week after week for a year. You can have big computers, but it’s the individuals who are devoted to the science that make the difference.”

Stress–Strain Can Build Strength

Whether studying an organic material, such as protein, or an inorganic material, such as metal, scientists must understand the origin of microstructural features to predict the material’s properties. In a metal, atoms stack in an orderly fashion, forming a crystalline lattice. However, some regions of the metal are less ordered. These disorderly areas have impurities or misaligned planes of atoms, called dislocations, that affect a metal’s properties.

Plasticity occurs in a metal as dislocations move, or slip, along well-defined crystal planes and directions. To predict a metal’s strength, researchers must understand its dislocation behavior. Dislocation dynamics—the interaction among dislocation lines—may be responsible for strain hardening, a property of metals in which a material’s strength increases as deformation increases.

The speed at which dislocations occur within an area in a metal can cause it to fail. An important part of ASC’s stockpile stewardship work is to examine metal strength under conditions of high strain rate and pressure. Materials scientist Vasily Bulatov leads a Grand Challenge team that is modeling in unprecedented detail the mechanics of dislocation motion by following the evolution of dislocation lines in response to applied stress. The team is using Thunder and the code Parallel Dislocation Simulator (ParaDiS) to develop an experimentally validated and



computationally efficient model of crystal strength. The team has focused on molybdenum and tantalum because these metals are similar to some materials in the nuclear stockpile. (See *S&TR*, November 2005, pp. 4–11.)

Dislocation dynamics simulations involve tracking millions of dislocation lines over millions of time steps. Thunder allows the team to surpass traditional computational studies that model systems two to three orders of magnitude smaller in size and shorter in time scale. “Thunder is the muscle ParaDiS needs,” says Bulatov. “We can investigate, for the first time, microstructural origins of stress–strain behavior.”

The team’s simulations are the first to show that during strain, colliding dislocations could partially merge into a junction bounded by two nodes. These multijunctions tie the dislocation lines into a tight knot that seems to confer strength in metals. The researchers also found that microstructures

with many multijunctions can multiply dislocations at a faster rate than microstructures with fewer multijunctions. The greater the number of multijunctions formed, the greater the amount of stress required to deform a metal (strain) before it breaks, and the harder the resulting metal.

Guiding Nanoscale Designs

The ability to predict material performance at micro- and nanoscales (billionths of a meter) helps researchers design chemical and biological detection instruments for homeland security. To guide the design of such small-scale instruments, Livermore researchers use simulations to predict how a component’s material properties may change as dimensions are reduced. Subtle changes in a material’s electronic and structural properties can be crucial to a microdevice’s performance. The simulations also predict how the device’s surrounding environment may change.

For example, in addition to the diverging theories on water's structural properties, debate exists regarding the properties of confined water, that is, water trapped between layers of another substance. Understanding how the structural and dynamic properties of water change as it passes through confined regions is extraordinarily challenging for both experimentalists and theorists. Physicist Eric Schwegler, who leads Livermore's Quantum Simulations Group, is conducting a Grand Challenge to simulate the electronic and structural properties of confined water at the nanoscale as it comes into contact with materials such as silicon carbide, graphite sheets, and carbon nanotubes. The team is focusing on these materials because their high electrical conductivity and biocompatibility are useful

characteristics for microfluidic devices designed to detect biothreat agents. (See *S&TR*, January/February 2006, pp. 11–17.)


To model the systems, the team uses a Livermore-developed, first-principles molecular dynamics code called Qbox. Designed for massively parallel systems, Qbox is used for conducting research on condensed matter subjected to extreme temperatures and pressures. The Qbox code calculates the electronic structure of a system with methods based on density functional theory, which defines a molecule's energy and electronic structure in terms of its electronic density.

The team requires about 560 of Thunder's processors and the Qbox code to simulate systems containing between 500 and 600 atoms for 20 to 25 picoseconds. Silicon and carbon have different electronic

properties, one of which is their affinity or aversion to water. One of the surprises the team found is that surfaces with an aversion to water tend to have a more pronounced effect on the structural and dynamical properties of confined liquid water than do surfaces with an affinity to water. For example, on the surface of graphite, which has an aversion to water, complex hydrogen-bond ring structures form at the surface. In confinement, the probability of forming these structures increases. Water molecules confined in carbon nanotubes demonstrate the same hydrophobic tendency.

The results can be used to prepare nanoscale materials that are most compatible for an intended purpose. They may also

Tom Arsenlis (left) and Vasily Bulatov (right) use the Parallel Dislocation Simulator code to study how metals deform and fail. The white areas in the inset show the formation of multijunctions, which tie three or more dislocations into tight knots, conferring strength to metal.




help protein-folding studies because the confined state of water may affect folding processes in an environment in which molecules have an aversion to water. "Access to Thunder through M&IC has enabled us to perform a predictive and systematic study of nanoscale confinement with respect to length scales, dimensionality, and interface effects without having to compromise on the level of theory," says Schwegler.

Schwegler's team includes collaborators from UC Berkeley, UC Davis, Massachusetts Institute of Technology, and Polytechnic of Torino, Italy. In fact, all of the Grand Challenge projects involve researchers outside the Laboratory. Because large volumes of data are produced in these studies, M&IC is developing the Green Data Oasis collaboration, which will enable external collaborators to more efficiently share data with Livermore researchers. Carnes says, "The powerful unclassified systems at the Laboratory produce vast amounts of numerical results. A storage capability outside the Laboratory's firewall allows our researchers to share data efficiently so they can achieve programmatic goals. The platform also enables UC students and faculty access to Livermore science."

Paving the Way for Atlas

When it isn't running simulations for one of the Grand Challenge projects, Thunder is being used to advance other scientific areas. For example, a team led by physicist Fred Streitz performed simulations for early studies on the nucleation and growth of grain structures in metals. The work was transferred to BlueGene/L, and in 2005, Streitz's team won the Gordon Bell Prize for reaching over 100 teraflops using a scientific application. (See *S&TR*, July/August 2006, pp. 17–19.) In another Livermore effort, physicist Evan Reed



A simulation of water molecules (red and white) confined within a carbon nanotube demonstrates carbon's hydrophobic tendency.

used Thunder to simulate coherent light from crystals. (See the article on p. 22.)

M&IC is currently gearing up for Atlas, a 44-teraflops machine that also will serve as the capability supercomputer for Grand Challenge projects beginning in January 2007. Because Atlas will roughly double Thunder's peak performance, Carnes expects Atlas simulations to make a large scientific impact on future Grand Challenge projects. "Consistently investing in institutional computing is an essential component of our science and technology investment strategy," says Carnes. "The addition of Atlas and other systems allows our researchers to push the limits of computing and its application in simulation science to ensure that Livermore remains a leader in science and technology."

—Gabriele Rennie

Key Words: 1BBL protein, Atlas supercomputer, climate change, confined water, dislocation dynamics, Green Data Oasis, Multiprogrammatic and Institutional Computing (M&IC) Program, Parallel Dislocation Simulator (ParaDiS) code, plasticity, protein folding, Qbox code, Replica Exchange Molecular Dynamics, Thunder supercomputer.

For further information contact Brian Carnes (925) 423-9181 (carnes1@llnl.gov).